

Simplicial homeomorphs and trace-bounded hypergraphs

Jason Long, Bhargav Narayanan, and Corrine Yap

ABSTRACT. Our first main result is the following basic fact about simplicial complexes: for each $k \in \mathbb{N}$, there exists an exponent $\lambda_k \geq k^{-2k^2}$ such that for any k -complex \mathcal{S} , every k -complex on $n \geq n_0(\mathcal{S})$ vertices with at least $n^{k+1-\lambda_k}$ facets contains a homeomorphic copy of \mathcal{S} . The existence of these exponents was suggested by Linial in 2006 but was previously known only in dimensions one and two, both by highly dimension-specific arguments: the existence of λ_1 is a result of Mader from 1967, and the existence of λ_2 was established by Keevash–Long–Narayanan–Scott in 2020. We deduce this geometric theorem from a purely combinatorial result about trace-bounded hypergraphs, where an r -partite r -graph H with partition classes V_1, V_2, \dots, V_r is said to be d -trace-bounded if for each $2 \leq i \leq r$, all the vertices of V_i have degree at most d in the trace of H on $V_1 \cup V_2 \cup \dots \cup V_i$. Our second main result is the following fact about degenerate trace-bounded hypergraphs: for all $r \geq 2$ and $d \in \mathbb{N}$, there exists an exponent $\alpha_{r,d} \geq (5rd)^{1-r}$ such that for any d -trace-bounded r -partite r -graph H , every r -graph on $n \geq n_0(H)$ vertices with at least $n^{r-\alpha_{r,d}}$ edges contains a copy of H . This strengthens a theorem of Conlon–Fox–Sudakov from 2009 who showed that a similar result holds for r -partite r -graphs H satisfying the stronger hypothesis that the vertex-degrees in all but one of its partition classes are bounded (in H , as opposed to in its traces).

1. INTRODUCTION

This paper aims to answer the following basic geometric question about k -dimensional simplicial complexes (or *k-complexes* for short) that arises in the ‘high-dimensional combinatorics’ programme of Linial [14, 15].

Problem 1.1. *Given a k -complex \mathcal{S} , how many facets can a k -complex on n vertices have if it contains no homeomorphic copy of \mathcal{S} ?*

For a k -complex \mathcal{S} , let $\lambda(\mathcal{S})$ be the supremum over all λ for which the maximum number of facets in a k -complex on n vertices with no homeomorphic copy of \mathcal{S} is $O(n^{k+1-\lambda})$. It is essentially folklore — see [9] for a discussion — that $\lambda(\mathcal{S}) > 0$ for every k -complex \mathcal{S} . A much more intriguing possibility, namely that for every k -complex

Date: 8 December, 2020.

2010 Mathematics Subject Classification. Primary 05E45; Secondary 05C65, 05C35.

\mathcal{S} , $\lambda(\mathcal{S})$ is bounded below uniformly by some universal exponent $\lambda_k > 0$ that depends only on the dimension k , was suggested by Linial [12, 13] (explicitly for dimension two and implicitly for higher dimensions); our first main result establishes this in every dimension.

Theorem 1.2. *For all $k \in \mathbb{N}$, there is a $\lambda_k \geq k^{-2k^2}$ such that for any k -complex \mathcal{S} , every k -complex on $n \geq n_0(\mathcal{S})$ vertices with at least $n^{k+1-\lambda_k}$ facets contains a homeomorphic copy of \mathcal{S} .*

The existence of such universal exponents λ_k as in Theorem 1.2 was previously known only for $k = 1$ and $k = 2$; that the optimal value of λ_1 is 1 is a classical result of Mader [16], and that $\lambda_2 \geq 1/5$ was shown recently by Keevash, Scott and the first and second authors [9]. The conjecturally optimal value of λ_2 is $1/2$, and establishing this remains open, though a beautiful recent result of Kupavskii, Polyanskii, Tomon, and Zakharov [11] establishes that $\lambda(\mathcal{S}) = 1/2$ whenever \mathcal{S} is the triangulation of any closed orientable two-dimensional surface, generalising a classical result of Brown, Erdős and Sós [2] establishing this fact for the two-sphere.

It is worth mentioning that all of [16, 2, 9, 11] proceed via arguments that are highly specific to dimensions one and two. Indeed, the main novelty in the proof of Theorem 1.2 is our ability to simultaneously handle all dimensions; this generality comes at a cost, however, since the resulting bounds in low dimensions are not very competitive with those in the aforementioned results.

We shall deduce Theorem 1.2 from a purely combinatorial result, of some independent interest, about the Turán numbers of a large class of r -uniform hypergraphs (or r -graphs, for short). For an r -partite r -graph H , let $\alpha(H)$ be the supremum over all α for which the maximum number of edges in an r -graph on n vertices with no copy of H is $O(n^{r-\alpha})$. A well-known result of Erdős [5] says that $\alpha(H) > 0$ for every r -partite r -graph H ; this value $\alpha(H)$ is called the *Turán exponent* of H , and the determination of these exponents is the central problem — see [7, 8] — of degenerate Turán theory.

To state our second result, we need a definition. We say that an r -partite r -graph H with partition classes V_1, V_2, \dots, V_r is d -trace-bounded if for each $2 \leq i \leq r$, all the vertices of V_i have degree at most d in the trace of H on $V_1 \cup V_2 \cup \dots \cup V_i$. Our second main result establishes the existence of universal lower bounds on the Turán exponents of degenerate trace-bounded hypergraphs.

Theorem 1.3. *For all $r \geq 2$ and $d \in \mathbb{N}$, there is an $\alpha_{r,d} \geq (5rd)^{1-r}$ such that for any d -trace-bounded r -partite r -graph H , every r -graph on $n \geq n_0(H)$ vertices with at least $n^{r-\alpha_{r,d}}$ edges contains a copy of H .*

Theorem 1.3 generalises a result of Conlon, Fox and Sudakov [3] which asserts, for all $r \geq 2$ and $d \in \mathbb{N}$, the existence of exponents $\alpha'_{r,d} > 0$ (of order roughly $(rd)^{1-r}$ as

well) with the following property: for any r -partite r -graph H with partition classes V_1, V_2, \dots, V_r in which the degrees of the vertices in each of V_2, V_3, \dots, V_r are at most d in H , we have $\alpha(H) \geq \alpha'_{r,d}$. It is easy to see that any r -partite r -graph H to which the aforementioned result applies is d -trace-bounded as well, so Theorem 1.3 clearly implies this result. Of course, Theorem 1.3 is genuinely stronger than the result in [3] since not every trace-bounded r -partite r -graph has bounded vertex-degrees in all but one of its partition classes, and indeed, the full strength of Theorem 1.3 will be crucial in proving Theorem 1.2.

Two more remarks about Theorem 1.3 are in order. First, the fact that the optimal value of $\alpha_{2,d}$ is $1/d$ (as opposed to the $1/(10d)$ promised by Theorem 1.3) is a result of Alon, Krivelevich and Sudakov [1] which may also be read out of earlier work of Füredi [6]. Second, it is known that, in a sense, something like the trace-boundedness hypothesis in Theorem 1.3 is necessary if one expects to control the Turán exponent in terms of vertex-degrees; indeed, from [10], we know that for every $\varepsilon > 0$, there exists a 3-partite 3-graph H with all the vertex-degrees in one of its partition classes being 1 for which $\alpha(H) \leq \varepsilon$.

Let us summarise the discussion above by specialising to 3-graphs. For a 3-partite 3-graph H with partition classes V_1, V_2 and V_3 , we have the following conclusions about its Turán exponent $\alpha(H)$, listed below in order of decreasing strength of the hypotheses on H .

- (1) If the degrees of the vertices in both V_3 and V_2 are bounded above by d in H , then [3] shows that $\alpha(H) \geq 1/(15d)^2$.
- (2) If the degrees of the vertices in V_3 are bounded above by d in H , and the degrees of the vertices in V_2 are bounded above by d in the trace of H on $V_1 \cup V_2$, then Theorem 1.3 says that $\alpha(H) \geq 1/(15d)^2$.
- (3) If all we know is that the degrees of the vertices in V_3 are bounded above by d in H , then [10] shows that $\alpha(H)$ need not be bounded below uniformly in terms of d , even when $d = 1$.

This paper is organised as follows. We begin by establishing some notation and making precise some of the undefined terminology appearing in the introduction in Section 2. The deduction of Theorem 1.2 from Theorem 1.3 is given in Section 3, and the proof of Theorem 1.3 follows in Section 4. We conclude with a discussion of some open problems in Section 5.

2. PRELIMINARIES

We shall only consider *homogeneous* k -complexes, namely k -complexes all of whose facets are k -dimensional. Consequently, we may specify a k -complex \mathcal{S} over a vertex set V by listing the family F of its k -dimensional facets, each of which is a subset of

V of cardinality $k + 1$ (though \mathcal{S} is, strictly speaking, the family of all subsets of its facets). We say that a k -complex \mathcal{T} contains a *homeomorph* (or a *homeomorphic copy*) of a k -complex \mathcal{S} if there is a subcomplex of \mathcal{T} that is homeomorphic to \mathcal{S} . An r -graph G on a vertex set V is a family E of r -element subsets of V called the edges of G . A k -complex \mathcal{S} may hence be identified with a $(k + 1)$ -graph G by viewing the facets of \mathcal{S} as the edges of G , and vice versa. When we specify a k -complex by its set of facets alone, or an r -graph by its edge set alone, the vertex set of the k -complex or the r -graph in question is taken to be the span of the facets or the edges respectively.

Since most of the work here will be in proving Theorem 1.3, let us set out some more notation for working with an r -graph G . We write $v(G)$ and $e(G)$ for the number of vertices and edges of G respectively. The *link* $\mathcal{L}(v, G)$ of a vertex $v \in V(G)$ in G is the $(r - 1)$ -graph whose edges are those sets S for which $\{v\} \cup S$ is an edge of G , and the *degree* $d(v, G)$ of v is the number of edges of G containing v , or equivalently $d(v, G) = e(\mathcal{L}(v, G))$. For an $(r - 1)$ -graph J with $V(J) \subset V(G)$, its *common neighbourhood* $\Gamma(J, G)$ in G is the set of vertices $v \in V(G)$ for which $\{v\} \cup S$ is an edge of G for each edge $S \in E(J)$. Finally, for a subset $U \subset V(G)$ of the vertex set of G , the *trace* $\text{Tr}(G, U)$ of G on U is the family $\{S \cap U : S \in E(G)\}$.

An r -graph G is *r-partite* if its vertex set admits a partition $V(G) = V_1 \cup V_2 \cup \dots \cup V_r$ such that every edge of G contains exactly one vertex each from each of these r partition classes. When G is an r -partite r -graph with partition classes V_1, V_2, \dots, V_r , we abbreviate $\text{Tr}(G, V_1 \cup V_2 \cup \dots \cup V_i)$ by $\text{Tr}_i(G)$, noting that $\text{Tr}_i(G)$ is an i -graph for each $1 \leq i \leq r$. Finally, we say that an r -partite r -graph G with partition classes V_1, V_2, \dots, V_r is *d-trace-bounded* if for each $2 \leq i \leq r$, we have $d(v, \text{Tr}_i(G)) \leq d$ for each $v \in V_i$.

It will be convenient for us to work with a large r -partite subgraph of a given r -graph; the following fact facilitates this, and follows from an easy averaging argument.

Proposition 2.1. *Any r -graph on rn vertices with m edges contains an r -partite subgraph with partition classes each of size n and at least $(r!/r^r)m$ edges.* \square

3. HOMEOMORPHS

Our proof of Theorem 1.2 relies on the following construction. Given a k -complex \mathcal{S} , the *canonical subdivision* of \mathcal{S} is a k -complex $\tilde{\mathcal{S}}$ homeomorphic to \mathcal{S} constructed as follows. The vertex set of $\tilde{\mathcal{S}}$ is given by

$$V(\tilde{\mathcal{S}}) = V(\mathcal{S}) \cup \{v_T : T \subset V(\mathcal{S}), |T| \geq 2, \text{ and } T \text{ is contained in some facet of } \mathcal{S}\};$$

in other words, we start with $V(\mathcal{S})$ and for each $2 \leq t \leq k + 1$, we introduce a new vertex v_T for each t -set T contained in some facet of \mathcal{S} . The facets $F(\tilde{\mathcal{S}})$ of $\tilde{\mathcal{S}}$ are then obtained by subdividing each facet of \mathcal{S} into $(k + 1)!$ facets as follows: for each facet




FIGURE 1. The construction of $\tilde{\mathcal{S}}$ for a 2-complex \mathcal{S} .

$S \in F(\mathcal{S})$ of \mathcal{S} , consider the $(k+1)!$ possible chains

$$\{v\} \subsetneq T_2 \subsetneq T_3 \subsetneq \cdots \subsetneq T_k \subsetneq S$$

with v a vertex of \mathcal{S} , and include $\{v, v_{T_2}, v_{T_3}, \dots, v_{T_k}, v_S\}$ in $F(\tilde{\mathcal{S}})$. It is not hard to see that $\tilde{\mathcal{S}}$ is homeomorphic to \mathcal{S} , as illustrated in Figure 1.

Proof of Theorem 1.2. We shall prove the result with $\lambda_k = \alpha_{k+1, (k+1)!}$, where $\alpha_{k+1, (k+1)!}$ is as promised by Theorem 1.3. We note that this establishes the bound

$$\lambda_k \geq (5(k+1)(k+1)!)^{-k} \geq k^{-2k^2}$$

for $k \geq 3$; that $\lambda_k \geq k^{-2k^2}$ for all $k \in \mathbb{N}$ follows from the facts, respectively from [16] and [9], that $\lambda_1 \geq 1$ and $\lambda_2 \geq 1/5$.

Given a k -complex \mathcal{S} , we first construct its canonical subdivision $\tilde{\mathcal{S}}$ as described above. When this k -complex $\tilde{\mathcal{S}}$ is viewed as a $(k+1)$ -graph, it is clear that it is $(k+1)$ -partite with partition classes V_1, V_2, \dots, V_{k+1} , where $V_1 = V(\mathcal{S})$ and for $2 \leq t \leq k+1$, V_t consists of those new vertices v_T introduced in $\tilde{\mathcal{S}}$ for each t -set T contained in some facet of \mathcal{S} . Furthermore, $\tilde{\mathcal{S}}$ is $((k+1)!)$ -trace-bounded; indeed, it is easy to see, for each $2 \leq t \leq k+1$, that for every $v \in V_t$, we have $d(v, \text{Tr}_t(\tilde{\mathcal{S}})) = t! \leq (k+1)!$.

It follows from Theorem 1.3 that provided $n \geq n_0(\mathcal{S})$ is large enough, any k -complex on n vertices with $n^{k+1-\lambda_k}$ facets contains $\tilde{\mathcal{S}}$ as a subcomplex, and therefore, a homeomorph of \mathcal{S} . \square

4. TRACE-BOUNDED HYPERGRAPHS

We start with a lemma that says that if an r -partite r -graph has many edges and a small number of (small) subgraphs that are ‘marked’ as being bad, then we may pass to an $(r-1)$ -partite $(r-1)$ -graph in one of its traces that has similar properties. To state this key lemma, we need a little set up, to which we now turn.

For each $r, d \in \mathbb{N}$, let $\mathcal{H}(r, d)$ denote the (finite) family of all nonempty r -partite r -graphs with at most d edges, taken up to isomorphism; recall our convention that the

vertex set of an r -graph specified by its edge set alone is the span of its edges, whence $r \leq v(J) \leq rd$ for each $J \in \mathcal{H}(r, d)$.

Suppose that $r \geq 2$, and let G be an r -partite r -graph with partition classes X_1, X_2, \dots, X_r each of size n . Suppose that for each $J \in \mathcal{H}(r, d)$, a subset $\mathcal{B}(J, G)$ of the copies of J in G have been *marked (as being bad)*. We say that an $(r-1)$ -graph $L \subset \text{Tr}_{r-1}(G)$ with at most d edges is β -bad with respect to G if

B1 either $|\Gamma(L, G)| \leq n^{1-\beta}$, or

B2 if there exists some $J \in \mathcal{H}(r, d)$ such that the number of marked copies of J in G containing L is at least

$$n^{-2\beta} (n^{1-\beta})^{v(J)-v(L)-1} |\Gamma(L, G)|.$$

The following lemma will be the workhorse that drives the proof of our main result.

Lemma 4.1. *For a fixed $r \geq 2$, $d \in \mathbb{N}$, and $\varepsilon, \delta > 0$, the following holds for all sufficiently large $n \in \mathbb{N}$. Let G be an r -partite r -graph with partition classes X_1, X_2, \dots, X_r each of size n and $e(G) \geq 2^r n^{r-\varepsilon}$ in which, for each $J \in \mathcal{H}(r, d)$, there is a set $\mathcal{B}(J, G)$ of at most $n^{v(J)-\delta}$ marked copies of J in G . Then, for $\beta = \delta/(rd+1)$, there is an $(r-1)$ -partite $(r-1)$ -graph $G' \subset \text{Tr}_{r-1}(G)$ with partition classes X_1, X_2, \dots, X_{r-1} such that*

- (1) $e(G') \geq 2^{r-1} n^{r-1-\varepsilon}$, and
- (2) for each $L \in \mathcal{H}(r-1, d)$, the set $\mathcal{B}(L, G')$ of copies of L in G' that are β -bad with respect to G satisfies

$$|\mathcal{B}(L, G')| \leq n^{v(L)-\beta+2\varepsilon}.$$

Proof. Choose a vertex $x \in X_r$ uniformly at random and let $G' = \mathcal{L}(x, G) \subset \text{Tr}_{r-1}(G)$ be the link of x in G . It is clear that

$$\mathbb{E}[e(G')] = e(G)/n \geq 2^r n^{r-1-\varepsilon}. \quad (1)$$

For each $L \in \mathcal{H}(r-1, d)$, let $P(L)$ be the set of copies L' of L in G' with $|\Gamma(L', G)| \leq n^{1-\beta}$. Next, for $L \in \mathcal{H}(r-1, d)$ and $J \in \mathcal{H}(r, d)$, we say that a copy J' of J in G *extends* a copy L' of L in $\text{Tr}_{r-1}(G)$ if $\text{Tr}_{r-1}(J') = L'$. Let $Q(L, J)$ be the set of copies L' of L in G' which extend to at least

$$n^{-2\beta} (n^{1-\beta})^{v(J)-v(L)-1} |\Gamma(L', G)|$$

marked copies of J belonging to $\mathcal{B}(J, G)$. With these definitions in place, we then have

$$\mathcal{B}(L, G') = P(L) \cup \left(\bigcup_{J \in \mathcal{H}(r, d)} Q(L, J) \right)$$

for each $L \in \mathcal{H}(r-1, d)$; indeed, the first term above accounts for **B1** and the second for **B2**.

First, we note that

$$\mathbb{E}[|P(L)|] \leq n^{v(L)-\beta}, \quad (2)$$

since each copy L' of L in $\text{Tr}_{r-1}(G)$ with $|\Gamma(L', G)| \leq n^{1-\beta}$ survives in G' with probability at most $n^{-\beta}$, and the number of copies of L in $\text{Tr}_{r-1}(G)$ is trivially at most $n^{v(L)}$.

Next, since $|\mathcal{B}(J, G)| \leq n^{v(J)-\delta}$ for each $J \in \mathcal{H}(r, d)$, we have

$$\sum_{L' \in Q(L, J)} n^{-2\beta} (n^{1-\beta})^{v(J)-v(L)-1} |\Gamma(L', G)| \leq |\mathcal{B}(J, G)| \leq n^{v(J)-\delta},$$

and by rearranging this, we get

$$\begin{aligned} \sum_{L' \in Q(L, J)} \frac{|\Gamma(L', G)|}{n} &\leq n^{2\beta} (n^{\beta-1})^{v(J)-v(L)-1} n^{v(J)-1-\delta} \\ &= n^{v(L)-\delta+\beta(1+v(J)-v(L))} \leq n^{v(L)-\delta+\beta rd}, \end{aligned}$$

where the last inequality uses the trivial facts that $v(J) \leq rd$ and $v(L) \geq 1$. It follows that

$$\mathbb{E}[|Q(L, J)|] \leq n^{v(L)-\delta+\beta rd} \quad (3)$$

for each $L \in \mathcal{H}(r-1, d)$ and $J \in \mathcal{H}(r, d)$.

Putting the estimates (2) and (3) together, we get

$$\mathbb{E}[|\mathcal{B}(L, G')|] \leq n^{v(L)-\beta} + |\mathcal{H}(r, d)| n^{v(L)-\delta+\beta rd} = C_1 n^{v(L)-\beta}, \quad (4)$$

where $C_1 = (1 + |\mathcal{H}(r, d)|)$, the last equality following from the fact that $\beta = \delta/(rd+1)$.

To finish, we set $C_2 = |\mathcal{H}(r-1, d)|$, and combine (1) and (4) to get

$$\mathbb{E} \left[\frac{e(G')}{2^{r-1} n^{r-1-\varepsilon}} - 1 - \frac{1}{C_2} \sum_{L \in \mathcal{H}(r-1, d)} \frac{|\mathcal{B}(L, G')|}{C_1 n^{v(L)-\beta}} \right] \geq 0;$$

consequently, there is at least one vertex in X_r whose link G' has the following properties:

- (1) $e(G') \geq 2^{r-1} n^{r-1-\varepsilon}$, and
- (2) for every $L \in \mathcal{H}(r-1, d)$, we have

$$|\mathcal{B}(L, G')| \leq \frac{C_1 C_2 n^{v(L)-\beta} e(G')}{2^{r-1} n^{r-1-\varepsilon}} \leq 2^{1-r} C_1 C_2 n^{v(L)-\beta+\varepsilon} \leq n^{v(L)-\beta+2\varepsilon},$$

here, we use the facts that $e(G') \leq n^{r-1}$, that $\varepsilon > 0$, and that n is sufficiently large.

Such an $(r-1)$ -graph G' has all the properties we require, proving the lemma. \square

With Lemma 4.1 in hand, we are now ready to prove our second main result.

Proof of Theorem 1.3. Let H be a d -trace-bounded r -partite r -graph with partition classes Y_1, Y_2, \dots, Y_r . For convenience, we prove that any large r -graph with sufficiently many edges on a vertex set whose cardinality is *divisible by r* must contain a copy of H ; of course, this divisibility assumption makes no material difference beyond allowing us to drop floor and ceiling signs. We shall prove the result with the precise value of

$$\alpha_{r,d} = \frac{1}{10d} \left(\frac{1}{rd+1} \right)^{r-2},$$

noting that $\alpha_{r,d} \geq (5rd)^{1-r}$ for all $r \geq 2$ and $d \geq 1$.

Given an r -graph on rn vertices with at least $(rn)^{r-\alpha_{r,d}}$ edges, then provided n is sufficiently large, we may, by Proposition 2.1, pass to an r -partite subgraph with partition classes X_1, X_2, \dots, X_r each of size n containing $2^r n^{r-\varepsilon}$ edges, for some

$$0 < \varepsilon \leq \frac{1}{9d} \left(\frac{1}{rd+1} \right)^{r-2};$$

we shall only work with this r -partite r -graph, which we call G , in what follows. Our goal now is to show that we are guaranteed to find a copy of H in G provided $n \geq n_0(H)$ is sufficiently large.

Our proof proceeds in two stages. In the first stage, we shall inductively construct a sequence of i -partite i -graphs $G_i \subset \text{Tr}_i(G)$ with partition classes X_1, X_2, \dots, X_i for $r-1 \geq i \geq 1$, with G_i being constructed by feeding G_{i+1} into Lemma 4.1. To accomplish this iterative construction, we need to find a suitable G_{r-1} from which to start, which we do as follows.

Claim 4.2. *There is an $(r-1)$ -partite $(r-1)$ -graph $G_{r-1} \subset \text{Tr}_{r-1}(G)$ with partition classes X_1, X_2, \dots, X_{r-1} such that*

- (1) $e(G_{r-1}) \geq 2^{r-1} n^{r-1-\varepsilon}$, and
- (2) *for each $J \in \mathcal{H}(r-1, d)$, the set $\mathcal{B}(J, G_{r-1})$ of copies of J in G_{r-1} that are contained in the link of fewer than $v(H)$ different vertices of X_r in G satisfies*

$$|\mathcal{B}(J, G_{r-1})| \leq n^{v(J)-1/2}.$$

Proof. The proof mirrors that of Lemma 4.1, but involves less work since we are aiming to accomplish less. Indeed, choose a vertex $x \in X_r$ uniformly at random and let $G' = \mathcal{L}(x, G) \subset \text{Tr}_{r-1}(G)$ be the link of x in G . As before, we clearly have

$$\mathbb{E}[e(G')] = e(G)/n = 2^r n^{r-1-\varepsilon}. \quad (5)$$

For each $J \in \mathcal{H}(r-1, d)$, the probability that a copy of J in $\text{Tr}_{r-1}(G)$ contained in the link of fewer than $v(H)$ different vertices of X_r survives in G' is at most $v(H)/n$, so it follows that $\mathcal{B}(J, G')$ satisfies

$$\mathbb{E}[|\mathcal{B}(J, G')|] \leq v(H) n^{v(J)-1}. \quad (6)$$

Setting $C = |\mathcal{H}(r-1, d)|v(H)$ and putting (5) and (6) together, we get

$$\mathbb{E} \left[\frac{e(G')}{2^{r-1}n^{r-1-\varepsilon}} - 1 - \frac{1}{C} \sum_{J \in \mathcal{H}(r-1, d)} \frac{|\mathcal{B}(J, G')|}{n^{v(J)-1}} \right] \geq 0.$$

Consequently, there is at least one vertex in X_r whose link G' has the following properties:

- (1) $e(G') \geq 2^{r-1}n^{r-1-\varepsilon}$, and
- (2) for every $J \in \mathcal{H}(r-1, d)$, we have

$$|\mathcal{B}(J, G')| \leq \frac{Cn^{v(J)-1}e(G')}{2^{r-1}n^{r-1-\varepsilon}} \leq 2^{1-r}Cn^{v(J)-1+\varepsilon} \leq n^{v(J)-1/2};$$

here, we use the facts that $e(G') \leq n^{r-1}$, that $\varepsilon < 1/2$, and that n is sufficiently large.

The claim follows by taking G_{r-1} to be a link G' with the above properties. \square

Let G_{r-1} be the $(r-1)$ -graph promised by Claim 4.2 and set $\delta_{r-1} = 1/2$. We know that

- (1) $e(G_{r-1}) \geq 2^{r-1}n^{r-1-\varepsilon}$, and
- (2) for each $J \in \mathcal{H}(r-1, d)$, the set $\mathcal{B}(J, G_{r-1})$ of copies of J in G_{r-1} contained in fewer than $v(H)$ distinct links in G satisfies $|\mathcal{B}(J, G_{r-1})| \leq n^{v(J)-\delta_{r-1}}$.

For $r-2 \geq i \geq 1$, having constructed G_{i+1} with

- (1) $e(G_{i+1}) \geq 2^{i+1}n^{i+1-\varepsilon}$ along with,
- (2) for each $J \in \mathcal{H}(i+1, d)$, a set $\mathcal{B}(J, G_{i+1})$ of at most $n^{v(J)-\delta_{i+1}}$ bad copies of J in G_{i+1} ,

we apply Lemma 4.1 to G_{i+1} to construct G_i such that

- (1) $e(G_i) \geq 2^i n^{i-\varepsilon}$, and
- (2) for each $J \in \mathcal{H}(i, d)$, the set $\mathcal{B}(J, G_i)$ of copies of J in G_i that are β_i -bad with respect to G_{i+1} satisfies

$$|\mathcal{B}(J, G_i)| \leq n^{v(J)-\delta_i},$$

where $\beta_i = \delta_{i+1}/((i+1)d+1)$ and $\delta_i = \beta_i - 2\varepsilon$.

Since $\delta_{r-1} = 1/2$ and $\delta_i = \delta_{i+1}/((i+1)d+1) - 2\varepsilon \geq \delta_{i+1}/(rd+1) - 2\varepsilon$ for each $r-2 \geq i \geq 1$, we get that

$$\delta_1 \geq \frac{1}{2} \left(\frac{1}{rd+1} \right)^{r-2} - 2\varepsilon \left(\frac{1 - (1/(rd+1))^{r-2}}{1 - 1/(rd+1)} \right) \geq \frac{1}{2} \left(\frac{1}{rd+1} \right)^{r-2} - 3\varepsilon > d\varepsilon,$$

where the last inequality above uses the fact that $\varepsilon \leq (1/9d)(1/(rd+1))^{r-2}$. Consequently, we have

- (1) $1/2 = \delta_{r-1} \geq \delta_{r-2} \geq \dots \geq \delta_1 > d\varepsilon$, and
- (2) $1/2 \geq \beta_{r-2} \geq \beta_{r-3} \geq \dots \geq \beta_1 > 0$.

We are now ready for the second stage of the proof where we embed H into G . First, we may assume that H has no isolated vertices, i.e., any vertices y for which $d(y, H) = 0$; indeed, if H has isolated vertices, we may embed the rest of H into G first, and then embed the isolated vertices into G arbitrarily provided n is sufficiently large. By assuming H has no isolated vertices, we have $\text{Tr}_r(H) = H$ and furthermore, the link of any vertex in $\text{Tr}_j(H)$ is nonempty for each $2 \leq j \leq r$.

For $1 \leq j \leq r-1$, we shall sequentially construct injective maps $\phi_j : Y_j \rightarrow X_j$ in such a way that

- (1) each induced map $\Phi_j : Y_1 \cup Y_2 \cup \dots \cup Y_j \rightarrow X_1 \cup X_2 \cup \dots \cup X_j$ is an embedding of $\text{Tr}_j(H)$ into G_j , and
- (2) for each nonempty subgraph J of $\text{Tr}_j(H)$ with at most d edges, no copy of J in the image of Φ_j is in $\mathcal{B}(J, G_j)$.

Observe that G_1 is a 1-graph on X_1 , which is just a subset of X_1 , so $v(G_1) = e(G_1)$, and we have

$$m = v(G_1) = e(G_1) \geq 2n^{1-\varepsilon}.$$

We construct $\phi_1 : Y_1 \rightarrow X_1$ in such a way that for every $J \in \mathcal{H}(1, d)$, no copy of J in $\phi_1(\text{Tr}_1(H))$ is in $\mathcal{B}(J, G_1)$, i.e., is β_1 -bad with respect to G_2 . By construction, we have $|\mathcal{B}(J, G_1)| \leq n^{v(J)-\delta_1}$ for each $J \in \mathcal{H}(1, d)$. Note that $\mathcal{H}(1, d)$ consists of d elements, namely one set of cardinality t for each $1 \leq t \leq d$, and also note that the number of β_1 -bad t -sets in G_1 is $o(m^t)$ for each $1 \leq t \leq d$ since $\delta_1/t \geq \delta_1/d > \varepsilon$. It follows that the number of problematic d -sets in G_1 , namely those containing a β_1 -bad t -set for some $1 \leq t \leq d$, is $o(m^d)$, so we may choose a subset of X_1 of size $|Y_1|$ which does not contain any such problematic d -set provided n is sufficiently large, as can be seen, for example, by applying a bound of de Caen [4] to the d -graph of all problematic d -sets. In other words, we can choose a subset S of X_1 of size $|Y_1|$ in such a way that for each $J \in \mathcal{H}(1, d)$, no copy of J in $\mathcal{B}(J, G_1)$ is contained in S ; we take ϕ_1 to be any injective map from Y_1 to S .

For $2 \leq j \leq r-2$, suppose that we have constructed injective maps $\phi_1, \phi_2, \dots, \phi_{j-1}$ as above. We now extend Φ_{j-1} to Φ_j by defining a suitable map $\phi_j : Y_j \rightarrow X_j$ randomly as follows.

Since H is d -trace-bounded, each vertex $y \in Y_j$ has degree at most d in $\text{Tr}_j(H)$. Given a vertex $y \in Y_j$, let $L(y) = \mathcal{L}(y, \text{Tr}_j(H))$ be the link of y in $\text{Tr}_j(H)$, so that $L(y)$ is a nonempty subgraph of $\text{Tr}_{j-1}(H)$ with at most d edges. Inductively, we know that $\Phi_{j-1}(L(y))$ is a subgraph of G_{j-1} . We choose $\phi_j(y)$ uniformly at random from the common neighbourhood $\Gamma(\Phi_{j-1}(L(y)), G_j) \subset X_j$. By choosing, for each $y \in Y_j$, the image $\phi_j(y)$ of y from the set $\Gamma(\Phi_{j-1}(L(y)), G_j)$, we have ensured that the induced map

Φ_j is a homomorphism from $\text{Tr}_j(H)$ into G_j . We claim that with positive probability, both of the following events hold:

- E1** the induced map Φ_j is injective, i.e., is an embedding of $\text{Tr}_j(H)$ into G_j , and
- E2** for each nonempty subgraph J of $\text{Tr}_j(H)$ with at most d edges, no copy of J in the image of Φ_j is in $\mathcal{B}(J, G_j)$.

To deal with **E1**, note that for any vertex $y \in Y_j$, the $(j-1)$ -graph $L(y) = \mathcal{L}(y, \text{Tr}_j(H))$ is nonempty and has at most d edges, and inductively, its image $\Phi_{j-1}(L(y))$ is not in $\mathcal{B}(L(y), G_{j-1})$, namely the set of copies of $L(y)$ in G_{j-1} that are β_{j-1} -bad with respect to G_j , so it follows that

$$|\Gamma(L(y), G_j)| \geq n^{1-\beta_{j-1}}. \quad (7)$$

Since $\beta_{j-1} \leq 1/2$, the probability that ϕ_j , and therefore Φ_j , fails to be injective is easily seen to be $o(1)$, whence we certainly have $\mathbb{P}(\mathbf{E1}) > 1/2$ provided n is sufficiently large.

To address **E2**, we argue as follows. Let J be a nonempty subgraph of $\text{Tr}_j(H)$ with at most d edges, and let $L = \text{Tr}_{j-1}(J)$. The probability of the event that $\Phi_j(J) \in \mathcal{B}(J, G_j)$ may be bounded above as follows. By the bound in (7), the number of choices for ϕ_j on $V(J) \setminus V(L)$ is at least $(n^{1-\beta_{j-1}})^{v(J)-v(L)}$. On the other hand, since L is nonempty with at most d edges, we know inductively that $\Phi_{j-1}(L) \notin \mathcal{B}(L, G_{j-1})$. Therefore, $\Phi_{j-1}(L)$ is not β_{j-1} -bad with respect to G_j , from which it follows that $\Phi_{j-1}(L)$ is contained in at most

$$n^{-2\beta_{j-1}} (n^{1-\beta_{j-1}})^{v(J)-v(L)-1} |\Gamma(\Phi_{j-1}(L), G_j)|$$

copies of J in G_j that belong to $\mathcal{B}(J, G_j)$. Hence, the probability of $\Phi_j(J) \in \mathcal{B}(J, G_j)$ is at most

$$\frac{n^{-2\beta_{j-1}} (n^{1-\beta_{j-1}})^{v(J)-v(L)-1} |\Gamma(\Phi_{j-1}(L), G_j)|}{(n^{1-\beta_{j-1}})^{v(J)-v(L)}} \leq \frac{n^{1-2\beta_{j-1}}}{n^{1-\beta_{j-1}}} = n^{-\beta_{j-1}} = o(1),$$

where we use the facts that $|\Gamma(\Phi_{j-1}(L), G_j)| \leq n$ and that $\beta_{j-1} > 0$. Summing this estimate over the $O(1)$ choices of nonempty subgraphs J of $\text{Tr}_j(H)$ with $e(J) \leq d$ shows that $\mathbb{P}(\mathbf{E2}) > 1/2$ provided n is sufficiently large as well, which together with the fact that $\mathbb{P}(\mathbf{E1}) > 1/2$ establishes the existence of an appropriate ϕ_j .

Now, we finish by extending Φ_{r-1} to an embedding $\Phi_r : Y_1 \cup Y_2 \cup \dots \cup Y_r \rightarrow X_1 \cup X_2 \cup \dots \cup X_r$ of H into G by defining a final injective map $\phi_r : Y_r \rightarrow X_r$. Recall that for each $J \in \text{Tr}_{r-1}(H)$, the set $\mathcal{B}(J, G_{r-1})$ consists of those copies J' of J in G_{r-1} for which the number of vertices $x \in X_r$ whose link $\mathcal{L}(x, G)$ contains J' as a subgraph is at most $v(H)$. We may now define ϕ_r by greedily picking, for each $y \in Y_r$, a distinct vertex $\phi_r(y) \in \Gamma(\Phi_{r-1}(\mathcal{L}(y, H)), G) \subset X_r$; this is always possible since $\Phi_{r-1}(\mathcal{L}(y, H)) \notin \mathcal{B}(\mathcal{L}(y, H), G_{r-1})$. It follows that $\Phi_r(H)$ is a copy of H in G , completing the proof. \square

5. CONCLUSION

A number of open problems remain, and we conclude by highlighting those that we think are particularly deserving of attention.

With regard to homeomorphs, the outstanding problem is to determine the optimal values of the exponents λ_k in Theorem 1.2 for each $k \geq 2$. As mentioned earlier, even the optimal value of λ_2 is not known, though there are good reasons to expect it to be $1/2$. In higher dimensions, we are only able to speculate: could it be that for each $k \geq 2$, the optimal value of λ_k is precisely $\lambda(S^k)$, where S^k is (any triangulation of) the k -sphere? This is indeed the underlying mechanism behind the prediction of the value of $1/2$ in dimension two, but we do not know, nor do we have a guess for, the value of $\lambda(S^k)$ for any $k \geq 3$; for example, a combination of a random construction and a generalisation of the argument of the Brown, Erdős and Sós [2] shows that $1/4 \leq \lambda(S^3) \leq 1/3$, but we have no reason to think either bound reflects the truth.

With regard to trace-bounded hypergraphs, the main problem again is to determine the optimal values of the exponents $\alpha_{r,d}$ in Theorem 1.3 for each $r \geq 3$ and $d \in \mathbb{N}$. As remarked upon earlier, we know that the optimal value of $\alpha_{2,d}$ is $1/d$, but even formulating a natural guess for the optimal value of $\alpha_{r,d}$ when $r \geq 3$, let alone proving it, would be of some interest.

ACKNOWLEDGEMENTS

The work of both the second author and the third author was supported in part by NSF grant DMS-1800521.

REFERENCES

1. N. Alon, M. Krivelevich, and B. Sudakov, *Turán numbers of bipartite graphs and related Ramsey-type questions*, Combin. Probab. Comput. **12** (2003), 477–494. 3
2. W. G. Brown, P. Erdős, and V. T. Sós, *On the existence of triangulated spheres in 3-graphs, and related problems*, Period. Math. Hungar. **3** (1973), 221–228. 2, 12
3. D. Conlon, J. Fox, and B. Sudakov, *Ramsey numbers of sparse hypergraphs*, Random Structures Algorithms **35** (2009), 1–14. 2, 3
4. D. de Caen, *Extension of a theorem of Moon and Moser on complete subgraphs*, Ars Combin. **16** (1983), 5–10. 10
5. P. Erdős, *On extremal problems of graphs and generalized graphs*, Israel J. Math. **2** (1964), 183–190. 2
6. Z. Füredi, *On a Turán type problem of Erdős*, Combinatorica **11** (1991), 75–79. 3
7. Z. Füredi and M. Simonovits, *The history of degenerate (bipartite) extremal graph problems*, Erdős centennial, Bolyai Soc. Math. Stud., vol. 25, János Bolyai Math. Soc., Budapest, 2013, pp. 169–264. 2

8. P. Keevash, *Hypergraph Turán problems*, Surveys in Combinatorics 2011, London Math. Soc. Lecture Note Ser., vol. 392, Cambridge Univ. Press, Cambridge, 2011, pp. 83–139. [2](#)
9. P. Keevash, J. Long, B. Narayanan, and A. Scott, *A universal exponent for homeomorphs*, Israel J. Math., To appear, arxiv:2004.02657. [1](#), [2](#), [5](#)
10. A. Kostochka, D. Mubayi, and J. Verstraëte, *Turán problems and shadows III: expansions of graphs*, SIAM J. Discrete Math. **29** (2015), 868–876. [3](#)
11. A. Kupavskii, A. Polyanskii, I. Tomon, and D. Zakharov, *The extremal number of surfaces*, Preprint, arxiv:2010.07191. [2](#)
12. N. Linial, *What is high-dimensional combinatorics?*, Random–Approx (2008). [2](#)
13. ———, *Challenges of high-dimensional combinatorics*, Lovász’s Seventieth Birthday Conference (2018). [2](#)
14. N. Linial and Z. Luria, *Discrepancy of high-dimensional permutations*, Discrete Analysis (2016:11), 8 pp. [1](#)
15. N. Linial and Y. Peled, *On the phase transition in random simplicial complexes*, Ann. of Math. **184** (2016), 745–773. [1](#)
16. W. Mader, *Homomorphieeigenschaften und mittlere Kantendichte von Graphen*, Math. Ann. **174** (1967), 265–268. [2](#), [5](#)

CAMBRIDGE CB1 1JH, UK

Email address: `jasonlong272@gmail.com`

DEPARTMENT OF MATHEMATICS, RUTGERS UNIVERSITY, PISCATAWAY, NJ 08854, USA

Email address: `narayanan@math.rutgers.edu`

DEPARTMENT OF MATHEMATICS, RUTGERS UNIVERSITY, PISCATAWAY, NJ 08854, USA

Email address: `corrine.yap@rutgers.edu`